CESifo WORKING PAPERS

7166
2018

July 2018

Structural Interpretation of Vector Autoregressions with Incomplete Information: Revisiting the Role of Oil Supply and Demand Shocks: Comment

Lutz Kilian, Xiaoqing Zhou

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7166*
Category 12: Empirical and Theoretical Methods

Structural Interpretation of Vector Autoregressions with Incomplete Information: Revisiting the Role of Oil Supply and Demand Shocks: Comment

Abstract

Recently, Baumeister and Hamilton (henceforth: BH) have argued that existing studies of the global oil market fail to account for uncertainty about their identifying assumptions. They recommend an alternative econometric approach intended to address this concern by formulating priors on the structural model parameters. We demonstrate that in practice BH are unable to parameterize identification uncertainty without falling back on ad hoc prior specifications. They are also unable to show that earlier studies did not impose all relevant identifying information. In fact, to the extent that BH's substantive conclusions differ from earlier studies, these differences do not reflect their use of a superior econometric methodology, but mainly the imposition of a highly unrealistic prior for the global impact price elasticity of oil supply. Once identification uncertainty about the global price elasticity of oil supply is accounted for by specifying a prior more in line with extraneous evidence and economic theory, the substantive results of earlier oil market studies are reaffirmed. We also refute BH's claim that existing oil market studies are invalid or not robust. Finally, we explain why the BH method is not a strict generalization of existing methods. It is, in fact, not designed to be applied to state-or-the-art oil market models because key assumptions of the proposed approach are not met in these models.

JEL-Codes: Q430, C320, E320.

Keywords: oil market models, structural VAR, identification, oil supply elasticity.

Lutz Kilian
Department of Economics
611 Tappan Street
USA – Ann Arbor, MI 48109-1220
lkilian@umich.edu

Xiaoqing Zhou
Bank of Canada
234 Wellington Street
Canada – Ottawa, ON, K1A 0G9
xzhou@bankofcanada.ca

July 15, 2018
We thank Ana Maria Herrera, Helmut Lütkepohl, Dan Murphy, and Pepe Montiel Olea for helpful comments. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Bank of Canada.

1. Introduction

Baumeister and Hamilton (2017) (henceforth: BH) argue that existing studies of the global oil market have failed to account for uncertainty about their identifying restrictions. Specifically, they single out three earlier oil market studies: One is the recursively identified oil market vector autoregressive (VAR) model of Kilian (2009), which started this literature. The other two studies are the sign-identified oil market VAR models of Kilian and Murphy (2012, 2014). BH argue that the prior distribution on the structural model parameters may be used to quantify the uncertainty about the identification restrictions, allowing this uncertainty to be captured by standard Bayesian estimation methods for structural VAR models (see Sims and Zha 1998). They insist that this “true Bayesian” approach is as a strict generalization of all existing frequentist and Bayesian methods for structural VAR analysis and allows them to relax the “strong” and “unreasonable” assumptions made in earlier work. BH suggest that using their approach reveals that oil supply shocks are a more important determinant of historical oil price fluctuations than in existing oil market models and that storage demand shocks are less important than previously thought. They also suggest that oil supply shocks are more recessionary for the global economy than shown by previous studies.

What are these conclusions based on? Perhaps surprisingly, BH do not actually examine whether allowing for identification uncertainty overturns the conclusions of earlier oil market studies such as Kilian (2009) or Kilian and Murphy (2012). Nor do they quantify the effect of allowing for identification uncertainty on the conclusions of Kilian and Murphy (2014). In fact, they do not present any quantitative analysis of the latter model at all. Instead, BH propose a new oil market VAR model that differs from existing models in multiple dimensions, and they evaluate this model based on their preferred econometric approach only. Their conclusions about the role of oil supply shocks are based on an informal comparison of the latter estimates with the estimates in Kilian and Murphy (2014) and related studies. Since the empirical results in question are derived using different model specifications, identifying assumptions, data sets, data transformations and estimators, however, the differences in the estimates cannot be attributed to identification uncertainty.

The actual explanation of their results is rather different. We stress that BH's preferred oil market VAR model has been specifically designed to inflate the effects of oil supply shocks and to reduce the effects of oil demand shocks. By far the most important tool used by BH to ensure results more favorable to oil supply shocks, however, is the imposition of a prior on the impact price elasticity of oil supply that attaches unrealistically large probability mass to high elasticity values. We show that their prior median for this elasticity is about 5 times larger than the largest credible microeconomic oil supply elasticity estimate and far from the benchmark of zero provided by economic theory. Under more economically plausible prior specifications, the substantive conclusions of earlier studies regarding the limited role of oil supply shocks in driving the real price of oil are reaffirmed (see Herrera and Rangaraju 2018).

2. Is the BH method a strict generalization of all existing econometric methods?

Consider a structural VAR process of order p,

$$B_0 y_t = B_1 y_{t-1} + ... + B_p y_{t-p} + w_t,\ t = 1,...,T,$$

with white noise innovations w_t, where the deterministic terms have been suppressed for simplicity. The structural shocks may be identified by a combination of normalization assumptions and identifying restrictions only on B_0^{-1}, only on B_0, or on both B_0^{-1} and B_0. Where the identifying restrictions are imposed depends on the economic context. In some cases, it is more natural to restrict the effect of structural shocks on observables, as summarized by B_0^{-1}, and in other cases the contemporaneous interaction of observables, as summarized by B_0 (see Kilian and Lütkepohl (2017). In the oil market literature, restrictions have typically been imposed on B_0^{-1}. In sign-identified structural oil market VAR models, in particular, it has been standard to impose restrictions on the signs of the elements of B_0^{-1} because typically economic models have direct implication for the sign of the responses of observables to structural shocks (see, e.g., Kilian and Murphy 2012, 2014; Antolin-Diaz and Rubio-Ramirez 2018).

The BH proposal is to impose a prior on individual elements of the matrix B_0 under the additional assumption that the marginal priors are independent across the elements of this matrix, which makes the method tractable. Effectively, BH expect researchers to change their economically motivated identifying restrictions, so the model can be fit into their restrictive framework. There are special cases in which the proposed method may be applied without substantive changes to the identifying assumptions. For example, when B_0^{-1} is recursively identified, as in Kilian (2009), so is B_0.

More generally, however, adapting existing nonrecursive oil market models to the BH methodology is difficult. A case in point is the oil market model of Kilian and Murphy (2012). BH are unable to implement this model, as originally proposed. The implementation of their method, as discussed in their paper, requires the imposition of additional exclusion restrictions on B_0 that are not required by Kilian and Murphy's (2012) model, that do not hold in the estimate reported by Kilian and Murphy (2012), and that are not implied by economic theory. Nor are BH able to handle the additional dynamic sign restrictions imposed on this model by Inoue and Kilian (2013).

In fact, applying the BH methodology to state-of-the-art oil market models such as the sign-identified model proposed by Kilian and Murphy (2014), which has also been utilized by Kilian and Lee (2014), Kilian (2017), Antolin-Diaz and Rubio-Ramirez (2018), and Herrera and Rangajaru (2018), among others, is impossible. One reason is that imposing priors on the price elasticity of oil demand requires restrictions on multiple elements of B_0, which violates the maintained assumption of independence in the BH method. This follows from the fact that the price elasticity of oil demand depends on the response of both oil inventories and oil production to the price change triggered by an exogenous oil supply shock. Another reason is that the BH approach, as implemented in their paper, does not allow for sign restrictions on impulse responses beyond the impact period.

Since their method cannot be applied to standard oil market models, it follows immediately that the BH approach is not a strict generalization of existing methods. It is a method designed for a subset of structural VAR models, which excludes many of the most interesting applications in applied work, and, even when it can be applied, it requires additional restrictive assumptions not employed by all existing studies. Thus, logic dictates that the proposed method is not a strict generalization of existing work.

BH emphasize that special cases of their method come close to reproducing the impulse response estimates in Kilian (2009) and Kilian and Murphy (2012). This fact does not imply that their analysis generalizes existing methods, however, as claimed by BH. Strictly speaking, they do not even generalize the analysis of the recursive model in Kilian (2009) since that analysis was frequentist, whereas the proposed new method is Bayesian. Even when the numerical estimates generated by Bayesian methods are numerically identical, their interpretation differs. Nor does the interpretation of Bayesian credible sets match that of frequentist confidence intervals. Thus, it would have been more correct for BH to state that their method allows them to incorporate the identifying restrictions used by Kilian (2009) within their specific Bayesian framework.

Another key difference between the BH approach and many earlier oil market VAR studies that complicates the comparison of the empirical results is the choice of the loss function underlying the impulse response estimators. Bayesian estimators are obtained from the posterior distribution of the statistic of interest by minimizing a given loss function in expectation. While BH claim that their Bayesian inference is optimal, it is optimal only conditional on their particular choice of the loss function. The specific loss function imposed by BH postulates that the user of the VAR model does not care about the shape and pattern of the response functions or the relationship across response functions. This assumption is economically unappealing because most researchers are very much interested in the joint evolution of the impulse responses. If we do not accept that loss function, the median response functions reported by BH are economically and statistically meaningless. They confound estimates from different structural models and, as has been demonstrated by example, they may severely distort the impulse response estimates (see, e.g., Fry and Pagan 2011, Kilian and Murphy 2012, Inoue and Kilian 2013, Kilian and Lütkepohl 2017). Although median responses functions may be numerically close to alternative estimators in some situations, in general, they cannot be expected to be the same. By restricting the loss function for the impulse responses, BH further reduce the generality of their approach.

3. Are the oil market models in Kilian (2009) and Kilian and Murphy (2012) invalid?

BH show that, after imposing the original identifying assumptions, their Bayesian impulse response estimates happen to lie almost on top of the least-squares estimates reported in Kilian (2009), confirming the robustness of the original conclusions under their alternative loss function. Much the same conclusion applies to BH's version of the Bayesian sign-identified VAR model of Kilian and Murphy (2012). What BH do not show is how imposing alternative priors on the identification affects these results, which is what one would have naturally expected them to show in their paper.

Instead, BH proceed by claiming that both the Kilian (2009) and the Kilian and Murphy (2012) model are fundamentally misspecified and by proposing their own oil market model. Their argument is that these earlier models, allegedly, imply large posterior probability mass on values of the short-run price elasticity of oil demand close to -2, which is much larger than any economically plausible elasticity value. BH attribute this result to the imposition of a short-run price elasticity of oil supply close to zero. They argue that the oil market models in Kilian (2009) and Kilian and Murphy (2012) should therefore be discarded.

As it turns out, this argument is wrong. The root cause is that, in producing their evidence, BH use an incorrect definition of the oil demand elasticity that ignores the storability of crude oil. BH's definition is based on the premise that at each point in time the production of oil and the consumption of oil are equal. As shown by Kilian and Murphy (2014), however, this premise is invalid, since oil is storable and oil inventory holdings change in response to oil supply shocks. In fact, the importance of accounting for storage in estimating the price elasticity of energy demand is widely recognized at this point (see, e.g., Coglianese, Davis, Kilian and Stock 2017). In short, the only reason that BH insist on using the wrong definition of the oil demand elasticity is that their econometric method, as discussed earlier, cannot handle oil market models subject to restrictions on the correctly defined elasticity.

Estimating the oil demand elasticity properly requires a structural model that includes changes in oil inventories. If one includes oil inventories in the oil market VAR model and makes the necessary modifications, as shown in Kilian and Murphy (2014), the incorrectly defined oil demand elasticity, as used by BH, drops from -2 to -0.44 and the correctly defined oil demand elasticity is only -0.25, not -2. This occurs in a model with a one-month oil supply elasticity close to zero, showing that the assumption of an inelastic short-run oil supply curve does not necessarily imply unrealistically large oil demand elasticities, as falsely implied by BH.

What happens when we drop oil inventories from the Kilian and Murphy (2014) oil market model and marginalize the time series process? This makes the coefficients in the 3-variable models of Kilian (2009) and Kilian and Murphy (2012) complicated transformations of the structural coefficients in the 4-variable Kilian and Murphy (2014) model, from which the oil demand elasticity value cannot be extracted (see Kilian and Lütkepohl 2017). As a result, the estimate near -2 reported by BH is not the impact price elasticity of oil demand at all. The latter elasticity is not identified in 3-variable oil market models. Kilian (2009) and Kilian and Murphy (2012) understood this point and hence did not discuss the estimation of the oil demand elasticity.

Another way of seeing that the "demand elasticity" issue is merely a smoke screen is to observe that the impulse responses and historical decompositions are remarkably robust as to whether a researcher works with the 3-variable or the 4-variable oil market framework.[1] Moreover, while BH cite the absence of restrictions on the oil demand elasticity in Kilian (2009) and Kilian and Murphy (2012) as an example of the "excessive confidence" of the authors, they fail to mention that the state-of-the-art oil market model of Kilian and Murphy (2014) already imposes economically motivated bounds on the oil demand elasticity. In fact, none of BH's concerns about the demand elasticity being unrealistically high apply in that model, yet the

[1] In fact, even the response estimates of the alternative model proposed by BH are much more similar to the estimates in Kilian (2009) and Kilian and Murphy (2012) – or for that matter in Kilian and Murphy (2014) – than to the Hamilton (2003) view that almost all oil price fluctuations are driven by oil supply shocks.

substantive conclusions of earlier oil market models such as Kilian (2009) and Kilian and Murphy (2012) remain largely unchanged, further illustrating that the concerns raised by BH are unwarranted.

4. What do we learn from the BH oil market model?

Given that BH are concerned with uncertainty about the identifying restrictions used in the literature, one has to wonder why they did not illustrate the alleged benefits of their econometric method based on Kilian and Murphy's (2014) oil market model, which after all is the most recent and most realistic model of its kind in the literature. The reason is simple. The challenge faced by BH is that their proposed econometric method cannot be applied to the Kilian and Murphy (2014) oil market model.

BH therefore design an oil market model of their own that is amenable to the application of their methodology and estimate that model on a different data set. Like the Kilian and Murphy (2014) model, the BH model includes oil inventories, but it does not impose the bounds on the properly defined impact price elasticity of oil demand, as discussed in section 2. In fact, this price elasticity of oil demand cannot even be computed based on the BH model. Rather BH impose their prior on the incorrect measure of the oil demand elasticity that equates oil production and consumption. BH then proceed to compare the estimates obtained from their model to the results in Kilian and Murphy (2014). The problem is that it is not clear what is driving the differences in the results because the BH oil market model also differs from the Kilian and Murphy (2014) model in many other dimensions. Specifically, differences may be due to (1) the choice of the loss function, (2) the model specification and the definition of the structural shocks, (3) the econometric methodology, (4) the choice of the prior, (5) the choice of the data and data transformations, (6) the number of autoregressive lags and the prior structure imposed on the lags, (7) the sample period, and (8) the definition of the oil demand elasticity.

4.1. Why the BH oil market model is questionable

What most of the modifications of the oil market model adopted by BH have in common is that they are designed to inflate the importance of oil supply shocks for the real price of oil and to dampen the estimated effect of oil demand shocks. For example, they substitute an arguably inferior measure of global real activity for the standard measure used in the earlier studies they seek to overturn (see Kilian and Zhou 2018). They then transform this alternative series to growth rates, which reduces the importance of flow demand shocks because it eliminates long cycles in the global business cycle.

Another good example is the selective introduction of measurement error. The apparent reason that BH choose to focus on measurement error in oil inventories exclusively is that they seek to lower the role for storage demand shocks and increase the role of oil supply shocks in their model. This approach is questionable. For example, why do BH only discuss measurement error in oil inventories, even though one can make as strong a case for measurement error in global oil production and even in the real price of oil (see Hamilton 2011)? Why do they do not mention that the conclusions in Kilian and Murphy (2014) regarding storage demand were validated against extraneous evidence?

The conventional approach of proxying for changes in global oil inventories based on OECD data has much appeal for most of the estimation sample. The single largest issue of measurement in oil inventories, as discussed in Kilian and Lee (2014), it is that the conventional oil inventory data fail to capture the rise in Chinese strategic oil inventories after 2010. Unfortunately, it is precisely this problem that the model of BH is not designed to handle. Not only can the systematic omission of Chinese oil inventories after 2010 not be captured by time-invariant additive mean zero measurement error in the level of the oil inventories, but BH's maintained assumption that OECD oil inventories are a time-invariant fraction of the unobserved true global oil inventories is incorrect by design. BH take the view that any arbitrary measurement error specification is better than none. It is not clear what the basis of that view is. It is also noteworthy that BH do not mention that Kilian and Lee (2014) reexamined the conclusions of Kilian and Murphy (2014) based on an alternative oil inventory series that includes Chinese inventories and obtained very similar results.

Another example of unwarranted changes in the analysis is the use of heavily regulated domestic oil price data for the United States from the pre-1973 era in estimating a global oil market model. BH argue that this fact constitutes a compelling advantage of their Bayesian approach. They ignore the fact that one cannot estimate a global oil market VAR model on data from an era when no such market existed. They also overlook the fact that oil price data from that era do not permit an autoregressive approximation (see, e.g., Alquist et al. 2013).

Yet another odd modeling choice is that BH impose a prior on the income elasticity of demand using industrial production as a proxy for national income. The problem is that industrial output accounts for only a fraction of national output (and hence national income) with a growing share of national income being derived from the output of the service sector, so the coefficient in question is not the income elasticity of oil demand (see Kilian and Zhou 2018).

4.2. What is really going on in the BH oil market model?

Although BH's alternative model replicates many of the key findings in the existing oil market literature about the role of oil demand shocks, BH do not dwell on this point, but instead stress that negative oil supply shocks in their alternative model are somewhat more powerful than in existing models such as Kilian and Murphy (2014). The problem from the point of view of the reader is that one cannot tell which of their modeling choices is responsible for this difference.

A recent study by Herrera and Rangaraju (2018) helps answer this question to some extent. They re-estimate both the Kilian and Murphy (2014) model and the BH model based on data for the same sample period. The Kilian and Murphy model is evaluated using the econometric methodology of Inoue and Kilian (2013, 2018) and the BH model based on their preferred econometric methodology. Herrera and Rangaraju show that the differences in the responses to oil supply shocks are heavily influenced by the prior that BH impose on the short-run price elasticity of oil supply. If that prior is replaced by a Student-t prior that is also truncated at the upper bound on the oil supply elasticity, as in Kilian and Murphy (2014), the BH model produces oil price responses to oil supply shocks that are as small as those in Kilian and Murphy (2014) and much smaller than reported by BH (see Figure 1). The peak response drops from 4.27 to 1.75. Similar results are also obtained for several larger upper bounds on the supply elasticity such 0.035 and 0.040. This does not mean that the impulse response dynamics are the same as in the Kilian-Murphy model, because the BH model was designed to inflate the responses to oil supply shocks in a variety of ways, but it shows that the findings in the existing literature about the limited role for oil supply shocks is extremely robust to changes in the model structure, econometric methods and data.

Which result we believe then largely depends on the realism of the prior for the impact price elasticity of oil supply. BH specify a truncated Student-t prior with 94% probability mass in excess of the upper bound on this elasticity imposed by Kilian and Murphy (2014) and with positive probability mass on elasticity values approaching infinity (see Figure 2). This approach is surprising given the overwhelming consensus in the literature that the short-run price elasticity is close to zero, which makes it natural to restrict the support of the prior.

The specific upper bound of 0.0258 on the global oil supply elasticity used in Kilian and Murphy (2014) was derived based on a simple thought experiment proposed by Kilian and Murphy (2012) at a time when microeconometric estimates of the short-run global price elasticity of oil supply were not available. More recent microeconomic estimates of the one-month oil supply elasticity are all close to zero and statistically insignificant. For example, Anderson, Kellogg and Salant (2018) report an estimate that is effectively zero for Texas oil producers. Related studies include Newell, Prest and Vissing (2016) and Bjørnland et al. (2017). Even though the latter study produced an elasticity estimate of 0.035, which is slightly larger than the upper bound in Kilian and Murphy (2012), that estimate is not statistically significant. Moreover, none of these studies refers to the global impact price elasticity of oil supply, which is what matters for the global oil market model. Rather they focus on regional data from the United States.[2] Thus, the upper bound in Kilian and Murphy (2012) is not invalidated by recent estimates of the one-month oil supply elasticity.

A tight upper bound is also consistent with economic theory. It is well established that changing oil production involves substantial adjustment costs. Anderson, Kellogg and Salant (2018) show that in the presence of such adjustment costs, oil producers should not adjust oil production in response to demand shifts, but should adjust their investment in future production, which implies an impact price elasticity of oil supply of zero effectively. The posterior median estimate in BH, in contrast, is 0.15, which makes it 4 times larger than even the largest impact price elasticity estimates from regional microeconomic data and inconsistent with economic theory (see Figure 2).

[2] The only study to focus on estimating the global price elasticity of oil supply is Bornstein, Krusell and Rebelo (2017). Their preferred IV estimate of the one-year oil supply elasticity is 0.12, which provides little information on the one-month global elasticity, since longer-run elasticities tend to be larger than short-run elasticities.

4.3. How do BH arrive at their oil supply elasticity estimate?

We can only speculate what prompted BH to choose their prior, but their choice can be understood in the context of the evolution of the literature on modeling oil prices. BH dispense with the upper bound on the oil supply elasticity imposed by earlier studies, allowing them to attach substantial probability weight to large values of the oil supply elasticity. BH refer to this as “relaxing” Kilian and Murphy’s upper bound. This language is deliberately misleading, because their approach really amounts to imposing positive probability mass on elasticity values larger than can be supported by extraneous evidence. Put differently, their results are not driven by the use of an informative prior as opposed to a degenerate or flat prior, as BH want us to believe, but by shifting this prior in the direction of unrealistically large supply elasticity values.

Why would BH do that? Hamilton for many years has been advocating the position that oil supply shocks have been the main driver of oil price fluctuations at least from the 1970s to the early 2000s, although the empirical support for this position was overturned many years ago (see, e.g., Kilian 2008a,b). Subsequently, Kilian and Murphy (2012) showed that a short-run price elasticity of oil supply close to zero rules out all models in which oil supply shocks have large effects on the real price of oil. In fact, even with a supply elasticity bound of 0.09, one can only explain 10% of the variability of oil prices. Thus, BH understood the importance of increasing the value of this elasticity in their effort to undermine existing evidence that oil demand shocks are the main driver of oil price fluctuations.

They accomplished this task mainly by imposing prior information on the distribution of the prior probability mass that reflects their personal beliefs rather than extraneous information. Since BH’s elasticity estimate is computed by integrating over the support of the posterior, the fact that they assign high posterior probability mass to ex ante implausible elasticity values pushes up the posterior estimate. If their prior had ruled out implausibly high elasticity values from the start, the authors would have obtained the same results as earlier studies. To put it bluntly, BH effectively impose the elasticity answer they want in an effort to inflate the oil price responses to oil supply shocks. What is deceptive about their study is that they attribute these differences to their econometric methodology, when in reality they simply impose these results in estimating the model.

4.4. Are there "alternative facts" about oil supply elasticities?

BH claim that the upper bound on the oil supply elasticity proposed by Kilian and Murphy (2012) is at odds with the facts. Of course, their claim that this upper bound is refuted by their own posterior median estimate of oil supply elasticity is circular, since that estimate was derived by putting high probability mass on unrealistically large elasticity values. They also argue that Kilian and Murphy (2012) in deriving this bound based on events in August 1990 failed to recognize that a threat by Saddam Hussein against the United Arab Emirates (UAE) in July 1990 caused the UAE to reduce its oil production in August. BH suggest that, adjusting for this change in oil production pushes the elasticity bound from 0.0258 up to 0.043. It is actually not clear at all whether this was the reason behind the production decline in the UAE in August because, according to the same sources, in July 1990 Saudi Arabia was pushing the UAE to cut back its production to push up the oil price, which provides an alternative explanation of this cutback.

One can certainly disagree on the motivations of oil producers in July 1990, but even the alternative upper bound of 0.043 is only one quarter as large as BH's preferred elasticity estimate of 0.15 and would immediately invalidate their posterior estimate. Moreover, Kilian and Murphy (2012) already showed that their own substantive findings are robust to any reasonable bound.[3] This result was recently reaffirmed by Herrera and Rangaraju (2018).

BH also claim to have alternative facts that show that the impact oil supply elasticity is much larger than 0.0258. For example, BH argue that oil supply disruptions associated with hurricanes, strikes in the oil industry, and attacks in oil fields can be used to estimate the impact price elasticity of oil supply. This argument does not make economic sense since we need an exogenous shift in the oil demand curve to identify the oil supply elasticity, not an exogenous shift in oil supply. It may seem that Kilian and Murphy's upper bound based on the response of global oil production to the real oil price increase caused by the invasion of Kuwait in August 1990 would be subject to the same critique. As discussed by Kilian and Murphy (2012), however, unlike in the examples cited by BH, a substantial part of the August 1990 increase in the real price of oil was caused by an exogenous shift in oil storage demand rather than an exogenous shift in oil supply. Moreover, to the extent that this price increase reflected the oil supply disruptions caused by the invasion, Kilian and Murphy stressed that this price increase was not only large, but it was expected to be persistent, and there was rare unanimity among oil producers about the need to respond as well as spare capacity, allowing us to view their response as an upper bound on what production response is feasible under favorable conditions. This makes the 1990 episode quite special and different from any of the oil supply disruptions cited by BH.

Nor does Figure 5 in the BH paper, which shows Saudi oil production from 1973 to 2014 with U.S. business cycle dates imposed, help pin down the impact price elasticity of oil supply. This plot is not informative about the global supply elasticity because (a) it relates to Saudi rather

[3] BH suggest that Kilian and Murphy (2012) themselves were uncertain about the upper bound because they reported alternative results with higher bounds. This is not the case. Kilian and Murphy (2012) designed the upper bound of 0.0258 to allow for reasonable doubt about higher elasticity values. The additional sensitivity analysis based on higher upper bounds was added to their paper in response to a request by a referee.

than global oil production, (b) what matters for the global oil market is exogenous shifts in global oil demand, not in U.S. demand, (c) we need, in addition, real oil price data for estimating the oil supply elasticity, and (d) all the relevant variables are fully endogenous, invalidating any causal interpretation of this figure. In short, it is not possible to infer the impact price elasticity of oil supply from this picture.

In addition, BH cite a drop in Saudi oil production between June 2008 and February 2009 as evidence of a higher one-month price elasticity of oil supply than assumed by Kilian and Murphy. This argument is misleading, because what we need for the model is the one-month supply elasticity, not some "elasticity" over eight months, because we need to control for the dynamic effects of earlier shocks, and because the reason why the price of oil fell over this period reflected both higher oil supply and lower oil lower demand (as shown in BH's Figure 10, for example). Again, the authors are trying to "estimate" the price elasticity of oil supply by looking at a picture.

Finally, although BH avoid citing the most relevant microeconomic studies on the price elasticity of oil supply, including studies by Anderson et al. (2018) and closely related work by Newell et al. (2016), which all find statistically insignificant supply elasticities close to zero, they do cite Bjørnland et al. (2017) for having estimated a one-month supply elasticity "as high as 0.2" for shale oil. The Bjørnland et al. supply elasticity estimate actually is 0.076 for shale oil and 0.035 for conventional oil (see their Table 2, last column). The 0.2 value (from the same Table cited by BH) is not the own-price elasticity of oil supply, as used in the VAR model, and hence is irrelevant.[4]

[4] The estimate of 0.2 referred to by BH is obtained by regressing the percent change in shale oil production on the oil futures spread. Unlike BH, Bjørnland et al. are careful not to refer to this coefficient estimate as a price elasticity of oil supply. The definition of the price elasticity of oil supply is the percent change in oil production induced by a one percent change in the real price of oil caused by an exogenous shift in oil demand. Thus, even if we interpret the

Not only is the Bjørnland et al. (2017) estimate of 0.035 not statistically significant, but, in addition, it should be noted that Bjørnland et al. are not estimating the global price elasticity of oil supply of interest for the VAR analysis, but the supply elasticity for North Dakota. Even if we take Bjørnland et al.'s North Dakota estimates as representative for the world (which is a big "if"), after taking account of the 4% share of shale oil in global oil production in 2014, the implied global oil supply elasticity would be only 0.96*0.035+0.04*0.076 = 0.037. Of course, shale oil production was not quantitatively important before 2011, as shown in Kilian (2017), so for much of the VAR estimation sample the correct computation is 0.035. Thus, the shale oil revolution had little effect on the impact price elasticity of oil supply in global oil markets. Either way, the elasticity is near 0.035, if we take the Bjørnland et al. estimates at face value.[5] The preferred estimate of 0.15 reported in BH is four times higher than this estimate and hence implausible. Thus, BH are unable to produce credible alternative evidence that supports their choice of the prior for the impact price elasticity of oil supply.

4.5. Are there new facts about oil demand elasticities?

BH argue at length that the impact price elasticity of oil demand cannot be -2, drawing on a wide range of estimates, some econometrically credible and some not, of the short-run and the long-run price elasticity of oil demand. There is no controversy about that fact. A quick look at the identifying assumptions in Kilian and Murphy (2014) should have sufficed to make this point, because Kilian and Murphy explicitly restricted the impact elasticity to be lower than the

---

oil futures spread as a proxy for the expected change in the spot price of oil, the regression coefficient in question is not the own-price elasticity of oil supply. In order to recover the shale oil supply elasticity from this regression, one would have to compute the change in the current real price of oil caused by the shift in inventory demand associated with an exogenous change in the expected oil price. Of course, expected oil price changes may not be exogenous, and, in the presence of a time-varying risk premium, the futures price will not be the expected change in the real price of oil, further complicating the analysis.

[5] This point is important because a shift in this elasticity would be incompatible with the maintained assumption of a time-invariant VAR model in the oil market literature. It would also be inconsistent with standard time-varying coefficient VAR models, because the latent parameters of this model would not be identified prior to 2011, when U.S. shale oil production became important (see Kilian 2017).

corresponding long-run elasticity of -0.8 suggested by microeconometric studies. Their posterior median estimate of this elasticity was -0.26, far from the bound of -0.8 and even further from the -2 value BH are concerned with. This point is important because it invalidates the claim by BH that previous work had ignored restrictions on the value of the oil demand elasticity and that BH bring new identifying information on this elasticity to bear. In short, BH in their paper claim credit for the idea of restricting the range of the oil demand elasticity, which was first proposed by Kilian and Murphy (2014) and is already standard in the literature.

4.6. How (not) to conduct sensitivity analysis

As noted earlier, the interpretation of the results in BH is complicated by the fact that they make numerous changes to existing oil market models. BH seek to address this concern by providing some sensitivity analysis. They assert that their results are robust to changes in the model assumptions, so identification uncertainty really is all that matters. This claim is deceptive. One obvious concern is that rather than removing their restrictive model features altogether, they merely place less weight on selected model features, resulting in models that are not very different from their baseline specification. A second concern is that some key model features are never relaxed at all. For example, no results are reported for 24 unrestricted autoregressive lags. Likewise, no results are reported for reasonable bounds on the oil supply elasticity or for VAR models without measurement error in oil inventories. A third concern is that BH change only one model feature at a time. They do not allow for the fact that the effect of relaxing these features need not be additive. For example, allowing for 24 rather than 12 autoregressive lags is unlikely to have much effect unless the real activity data are transformed as in earlier studies and the tight prior on the lag structure is relaxed at the same time. Thus, BH's sensitivity analysis obscures more than it illuminates what is driving their estimates.

5. BH's critique of the Kilian and Murphy (2014) oil market model

BH go out of their way not to evaluate the Kilian and Murphy (2014) model on its own terms. As we already discussed, the reason is that their proposed method cannot handle this model. Instead, they devote a subsection to disparaging this model. It is important to clarify this discussion. First, BH claim that the results based on this model cannot be replicated or updated. This claim is false. Successful replications and updates include Kilian and Lee (2014), Kilian (2017), Antolin-Diaz and Rubio-Ramirez (2018), and Herrera and Rangaraju (2018), among others. In fact, the original methodology in Kilian and Murphy (2014) has been subsequently refined to account for estimation uncertainty, using state-of-the-art methods that dispense with the economically implausible loss function imposed by BH. Efficient codes for estimating oil market VAR models have been readily available for many years and are routinely used in the literature. Nor is the fraction of surviving draws an issue in this class of models. Merely changing the solution algorithm can easily increase the number of draws by orders of magnitude (see Kilian and Lütkepohl 2017). Finally, the primary reason that BH obtain many more posterior draws from their econometric model is not the superiority of their method, but that they fail to impose the oil supply elasticity bound and other identifying restrictions.

Second, it is worth mentioning that Kilian and Murphy (2014) verified that the model they selected among the set of admissible models matches extraneous evidence about the timing of shifts in inventory demand.[6] In Figure 8 of their paper, BH present an alternative model solution obtained by selecting another random seed, which is intended to illustrate the lack of robustness of the impulse response estimates in Kilian and Murphy (2014). Unlike Kilian and Murphy (2014), however, BH do not appear to conduct an external validation of this alternative model solution. Thus, their analysis is incomplete. Moreover, if BH had been serious about

[6] For further discussion of this point see also Kilian and Lee (2014) and Kilian and Lütkepohl (2017).

comparing these model solutions, they would have compared their posterior density values, as proposed in Inoue and Kilian (2013, 2018). Figure 3 illustrates that the responses of the most likely model solution, in fact, look very much like those of the model selected by Kilian and Murphy (2014), even when the sample is updated to August 2017, directly refuting BH's claim that the estimates in Kilian and Murphy (2014) are not robust and cannot be updated.

Third, BH suggest that the measure of global real economic activity employed by Kilian (2009) and Kilian and Murphy (2012, 2014) is "unconventional". This claim is surprising since the Kilian index is arguably the most widely used index of its kind in the literature. BH also argue that the Kilian (2009) global real economic activity index has no clear economic interpretation. The specific claims made by BH in this context have been shown to be false and confused in recent work by Kilian and Zhou (2018) who also discuss problems with the use of the proxy for global industrial production favored by BH. It is worth adding that, contrary to the statements by BH, an industrial production index is an index and hence does not have natural units (such as tons or bushels) nor is a lack of natural units a problem for estimating or interpreting oil market models.

Fourth, there is no compelling reason to replace the oil inventory specification in Kilian and Murphy (2014). BH do not establish the existence of measurement error in oil inventories. Rather they impose a particularly unappealing form of measurement error on this variable. Finally, there are compelling reasons to ignore the oil market data prior to 1973, as is standard in the literature. The inclusion of pre-1973 data in the estimation, as favored by BH, does not make economic sense, given the discrete-continuous time series process of the price of oil prior to 1973 and given the well-documented structural break in the data in late 1973 (see Dvir and Rogoff 2009; Alquist et al. 2013).

6. Is allowing for identification uncertainty necessarily a good idea?

Even granting that, in principle, accounting for identification uncertainty would be attractive, the BH paper illustrates that it is difficult to describe and parameterize this uncertainty in practice. The premise in BH that researchers have failed to utilize extraneous information about elasticities is clearly incorrect. Although BH stress the benefits of using such information, they provide no guidance as to how to map extraneous evidence into a prior distribution. For example, a few more or less relevant microeconometric point estimates of elasticities are not enough to form a prior. Nor do BH explain where to obtain such additional prior information. Their own priors are driven by computational tractability and convenience rather than realism. There is no question that imposing ad hoc priors on parameters of interest may do more harm than good in analyzing oil market VAR models. The truncated Student-t prior favored by BH is a case in point.

In the case of the impact price elasticity of oil supply, for example, values near zero are economically more plausible than larger values, suggesting a truncated prior distribution for the supply elasticity that reaches its maximum at zero (the benchmark provided by economic theory) and smoothly decays to zero at the upper bound. Relative to this benchmark, the concern with the truncated uniform prior implicitly imposed by Kilian and Murphy, if any, is that it imposes too much probability weight on elasticities close to the upper bound, which biases the results away from low supply elasticity values. Thus, what BH should have done is to compare the effects of such an informative elasticity prior with that of a uniform prior. Given the results in Herrera and Rangaraju (2018), a fair conjecture is that BH would have concluded that accounting for identification uncertainty is of lesser importance in this case.

Moreover, while such an economically informed prior for the oil supply elasticity would have been a step in the right direction, it would not have been enough to resuscitate the BH approach. The problem remains that the BH approach does not allow us to impose a prior on the properly defined impact price elasticity of oil demand. Thus, the problem of dealing with uncertainty about identifying restrictions remains unresolved.

7. Conclusions

BH make far-reaching claims about (1) the invalidity of existing oil market models, (2) the invalidity of existing econometric methods, (3) the benefits and generality of their proposed alternative econometric method, and (4) the determinants of the real price of oil. We demonstrated that none of these claims is supported by their evidence. Specifically, we identified six main areas of concern with their analysis. First, the BH method is not a strict generalization of existing methods. It is in fact not designed to be applied to state-or-the-art oil market models because key assumptions of this econometric method are not met in these models. Second, BH fail to demonstrate how identification uncertainty can be parameterized in practice without falling back on ad hoc prior specifications. Third, they are also unable to show that earlier studies did not impose all relevant identifying information. Fourth, to the extent that BH's substantive conclusions differ from earlier studies, these differences do not reflect their use of a superior econometric methodology, but mainly the imposition of a highly unrealistic prior for the impact price elasticity of oil supply. Once identification uncertainty about the price elasticity of oil supply is accounted for by specifying a prior more in line with extraneous evidence and economic theory, the substantive results of earlier oil market studies are reaffirmed (see Herrera and Rangaraju 2018). Fifth, BH's claim that existing oil market models imply unrealistically large impact price elasticities of oil demand and hence should be discarded is false, as is their claim that the results in Kilian and Murphy (2014) are not robust. Sixth, BH purposely modified the standard oil market model in order to reduce the effects of oil demand shocks and to inflate the effects of oil supply shocks. BH do not control for these and other questionable changes in the analysis and hence incorrectly attribute the differences in the results to identification uncertainty. Finally, to the extent that BH provide sensitivity analysis, they do so only selectively and incompletely, and they do not consider the combined impact of all the changes they make to conventional oil market models, making it difficult to interpret their results.

References

Alquist, R., Kilian, L., and R.J. Vigfusson (2013), "Forecasting the Price of Oil," in: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland, 2013, 427-507.

Anderson, S.T., Kellogg, R., and S.W. Salant (2018), "Hotelling Under Pressure," *Journal of Political Economy*, 126, 984-1026.

Antolin-Diaz, J., and J.F. Rubio-Ramirez (2018), "Narrative Sign Restrictions," forthcoming: *American Economic Review*.

Baumeister, C., and J.D. Hamilton (2017), "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Demand and Supply Shocks," manuscript, University of California at San Diego.

Bjørnland, H.C, Nordvik, F.M., and M. Rohrer (2017), "Supply Flexibility in the Shale Patch: Evidence from North Dakota," manuscript, Norwegian Business School.

Bornstein, G., Krusell, P., and S. Rebelo (2017), "Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry," manuscript, Institute for International Economic Studies.

Coglianese, J., Davis, L.W., Kilian, L., and J.H. Stock. (2017), "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand," *Journal of Applied Econometrics*, 32, 1-15.

Dvir, E., and K. Rogoff (2009), "Three Epochs of Oil," NBER Working Paper No. 14927.

Fry, R., and A.R. Pagan (2011), "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 49, 938-960.

Hamilton, J.D. (2003), "What Is an Oil Shock?" *Journal of Econometrics*, 113, 363-398.

Hamilton, J.D. (2011), "Nonlinearities and the Macroeconomic Effects of Oil Prices," *Macroeconomic Dynamics*, 15, 364-378.

Herrera, A.M., and S.K. Rangaraju (2018), "The Effect of Oil Supply Shocks on U.S. Economic Activity: What Have We Learned?", manuscript, University of Kentucky.

Inoue, A., and L. Kilian (2013), "Inference on Impulse Response Functions in Structural VAR Models," *Journal of Econometrics*, 177, 1-13.

Inoue, A. and L. Kilian (2018), "Corrigendum to 'Inference on impulse response functions in structural VAR models' [J. Econometrics 177 (2013), 1-13]," (with A. Inoue), forthcoming: *Journal of Econometrics*.

Kilian, L. (2008a), "The Economic Effects of Energy Price Shocks," *Journal of Economic Literature*, 46, 871-909.

Kilian, L. (2008b), "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics*, 90, 216-240.

Kilian, L. (2009), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market", *American Economic Review*, 99, 1053-1069.

Kilian, L. (2017), "The Impact of the Fracking Boom on Arab Oil Producers," *Energy Journal*, 38(6), 137-160.

Kilian, L., and T.K. Lee (2014), "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories," *Journal of International Money and Finance*, 42, 71-87

Kilian, L., and H. Lütkepohl (2017), *Structural Vector Autoregressive Analysis*, Cambridge: Cambridge University Press.

Kilian, L., and D.P. Murphy (2012), "Why Agnostic Sign Restrictions Are Not Enough:

Understanding the Dynamics of Oil Market VAR Models," *Journal of the European Economic Association*, 10, 1166-1188.

Kilian, L., and D.P. Murphy (2014), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*, 29, 454-478.

29, 454-478.

Kilian, L., and T.K. Lee (2014), "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories," *Journal of International Money and Finance*, 42, 71-87.

Kilian, L., and X. Zhou (2018), "Modeling Fluctuations in the Global Demand for Commodities," forthcoming: *Journal of International Money and Finance*.

Newell, R., Prest, B., and A. Vissing (2016), "Trophy Hunting vs. Manufacturing Energy: The Price Responsiveness of Shale Gas," manuscript, Resources for the Future.

Sims, C.A., and T. Zha (1998), "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39, 949-968.

Figure 1: How Imposing the Elasticity Bound Reduces the Response of the Real Price of Oil to an Unexpected Oil Supply Disruption

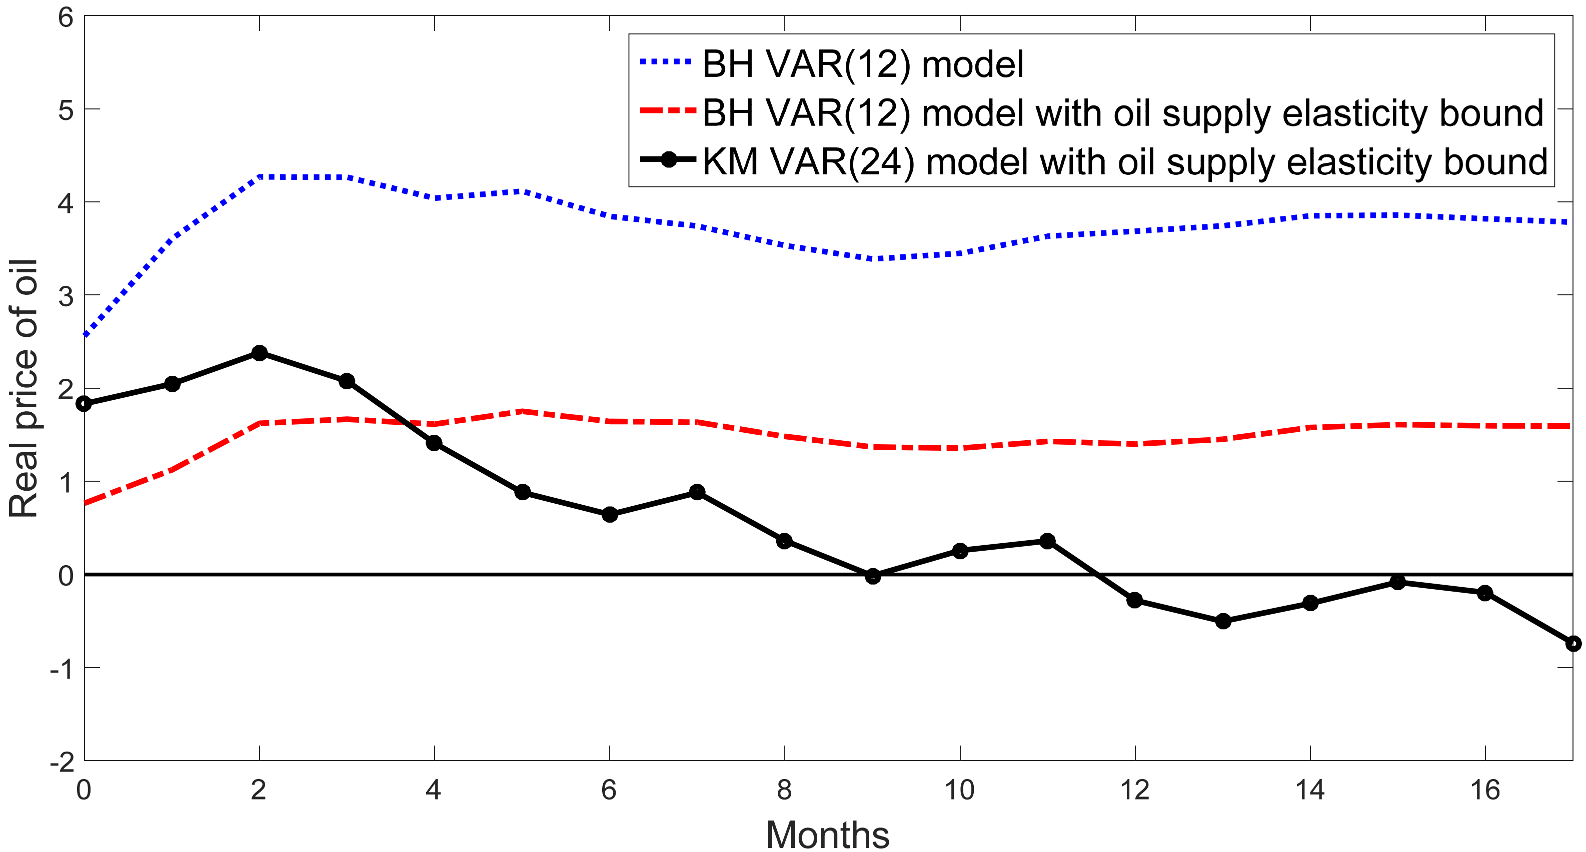

NOTES: These estimates are courtesy of Herrera and Rangaraju (2018). The BH model estimates refer to pointwise posterior median responses. The estimate based on the updated Kilian and Murphy (2014, KM) model is the response function of the modal admissible structural model (see Inoue and Kilian 2013, 2018). The BH and KM models and estimators differ in many dimensions, but the estimates agree that the oil price response to oil supply shocks is small compared with the response to oil demand shocks, once the impact price elasticity of oil supply is bounded by [0,0.0258].

Figure 2: BH's prior for the one-month global price elasticity of oil supply

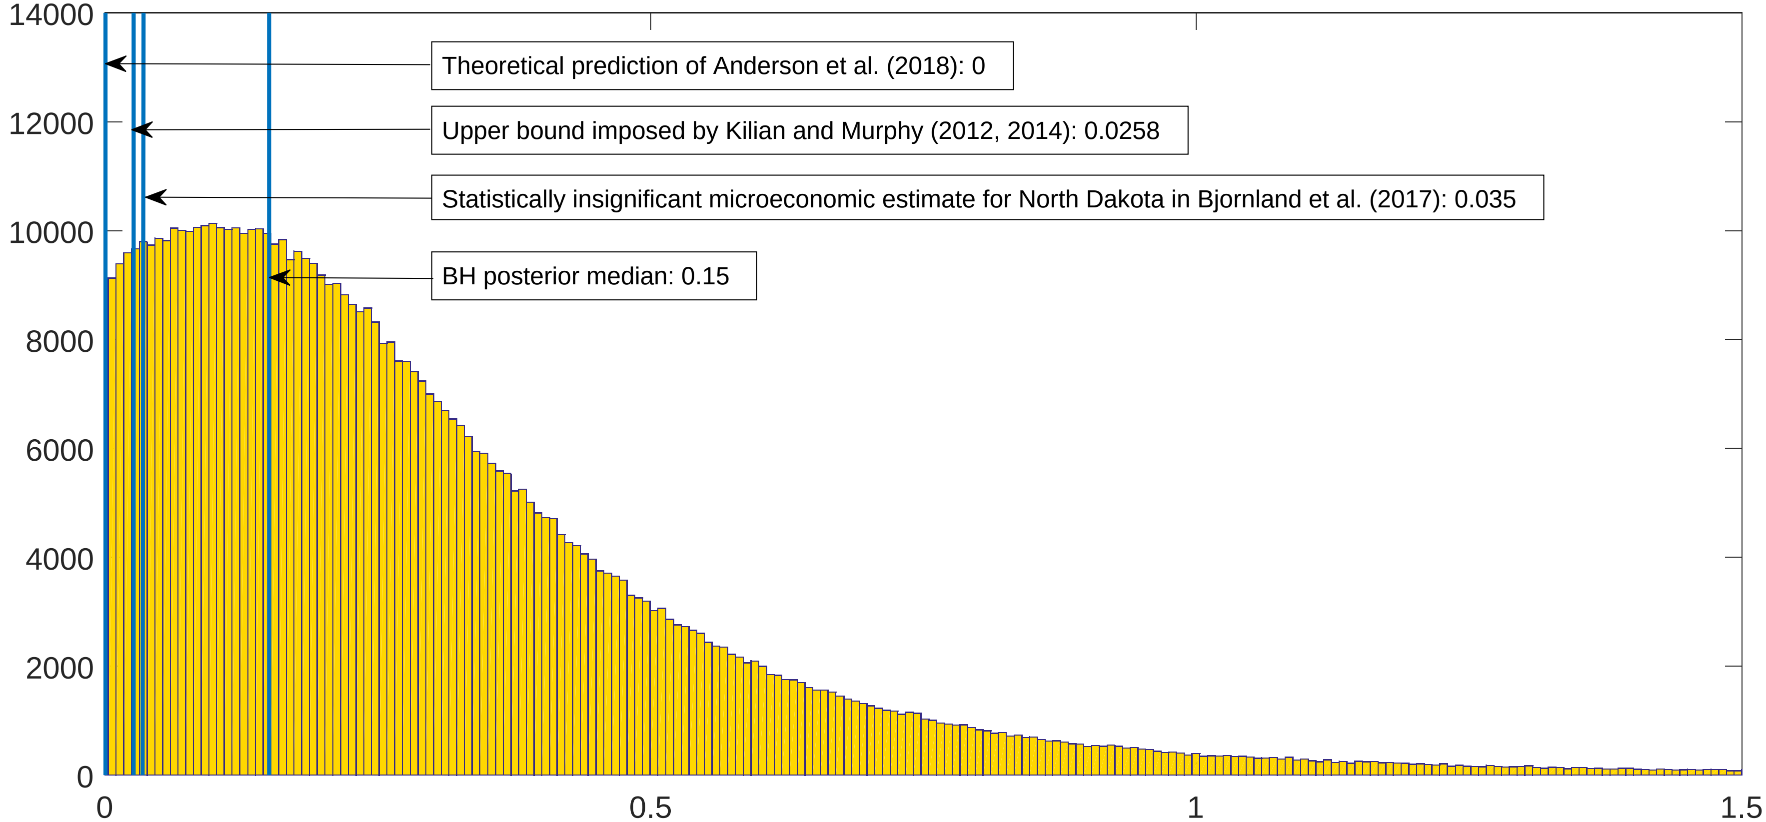

NOTES: The prior was constructed as described in Table 1 of BH. BH assign 94% prior probability mass to elasticity values larger than the upper bound in Kilian and Murphy (2012 2014). The BH prior density is unbounded to the right and assign positive probability mass to supply elasticities approaching infinity. The prior median is five times larger than the estimate in Bjørnland et al. (2017) and the implied posterior median is four times larger.

Figure 3: The robustness of the Kilian and Murphy (2014) response estimates

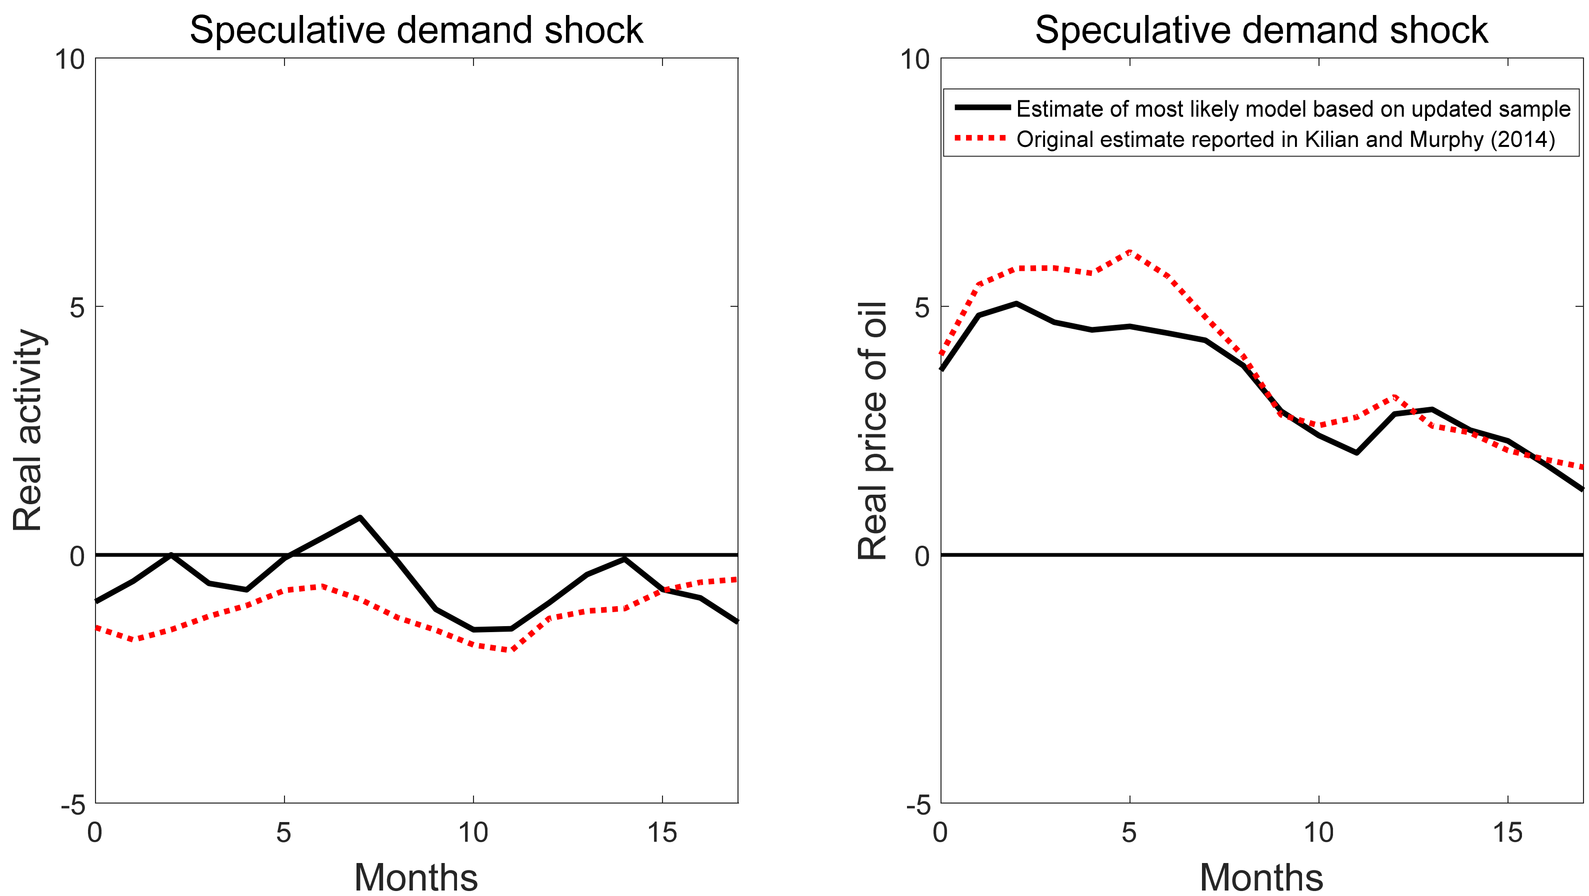

NOTES: The plot shows the responses of the original estimate selected by Kilian and Murphy (2014) as well as the corresponding responses from the most likely model obtained when evaluating the Kilian and Murphy model on updated data ending in August 2017. The most likely model was selected using the methodology of Inoue and Kilian (2013, 2018). The two estimates are quite similar, contrary to the claim made by BH.